

EXAMINER'S AMENDMENT

An examiner's amendment to the record appears below. Should the changes and/or additions be unacceptable to applicant, an amendment may be filed as provided by 37 CFR 1.312. To ensure consideration of such an amendment, it MUST be submitted no later than the payment of the issue fee.

Authorization for this examiner's amendment was given in a telephone interview with Mr. Mr. Noel Kivlin (Reg. No.: 33929) on Wednesday, October 08, 2008.

The application has been amended as follows:

In the claims:

*** Please amend/replace all claims as follows:**

1-31. (Cancelled)

32. (Currently Amended) A computer-readable storage medium storing program instructions executable to implement a method comprising:

storing a respective replica of an object on each respective node of a plurality of nodes; a first node of the plurality of nodes initiating communication a first distributed transaction with each of the other nodes of the plurality of nodes to attempt to synchronously update the replicas of the object stored on each of the other nodes of the plurality of nodes, wherein the communication is successful for each respective node of a first subset of the other nodes successfully commits the first distributed transaction to update the replica of the object stored on the respective node, and unsuccessful for wherein each node of a second subset of the other nodes does not successfully commit the first distributed transaction; and

for each respective node of the first subset of the other nodes;

the respective node updating the replica of the object stored on the respective node;
and

each the respective node of the first subset of the other nodes adding an identification of the object to a respective list of incoherent objects stored on the respective node in response to the communication being unsuccessful for the second subset of the other nodes determining that not all of the other nodes successfully committed the first distributed transaction;

wherein the first subset of the other nodes includes a second node of the plurality of nodes; wherein the method implemented by the program instructions further comprises:

after the second node adds the identification of the object to the respective list of incoherent objects stored on the second node, the second node attempting to communicate with each of the plurality of nodes other than the second node; and

in response to successfully communicating with each of the plurality of nodes other than the second node, the second node initiating a second distributed transaction to synchronize the replicas stored on the second subset of the other nodes with the replicas stored on the first subset of the other nodes.

33. (Currently Amended) The computer-readable storage medium of claim 32, wherein the method implemented by the program instructions further comprises:

the first node updating the replica of the object stored on the first node; and

the first node adding an identification of the object to a respective list of incoherent objects stored on the first node in response to the communication being unsuccessful for the second subset of the other nodes determining that not all of the other nodes successfully committed the first distributed transaction.

34. (Canceled)

35. (Currently Amended) The computer-readable storage medium of claim [[34]] 32, wherein the first subset of the other nodes includes at least a quorum of the plurality of nodes, successfully commit the transaction.

36. (Canceled)

37. (Currently Amended) The computer-readable storage medium of claim [[36]] 32,

wherein the method implemented by the program instructions further comprises each respective node of the first subset of the other nodes removing the identification of the object from the respective list of incoherent objects stored on the respective node in response to said synchronizing the replicas stored on the second subset of the other nodes with the replicas stored on the first subset of the other nodes.

38. (Currently Amended) The computer-readable storage medium of claim [[36]] 32,
wherein each respective node of the first subset of the other nodes updates the replica of the object stored on the respective node by applying a first change to the replica of the object stored on the respective node;

wherein synchronizing the replicas stored on the second subset of the other nodes with the replicas stored on the first subset of the other nodes comprises each respective node of the second subset of the other nodes applying the first change to the replica of the object stored on the respective node.

39. (Currently Amended) The computer-readable storage medium of claim 32,
wherein the object is a first object of a plurality of objects, wherein each object of the plurality of objects has a plurality of replicas stored on the plurality of nodes;
~~wherein the first subset of the other nodes includes a second node;~~
wherein the respective list of incoherent objects stored on the second node includes identifications of ~~two or more one or more additional objects~~ of the plurality of objects ~~in addition to the first object;~~

wherein the method implemented by the program instructions further comprises:
for each respective object ~~of the one or more additional objects~~ identified in the list of incoherent objects stored on the second node, the second node attempting to communicate with the plurality of replicas of the respective object to initiate a ~~respective~~ distributed transaction to synchronize the plurality of replicas of the respective object.

40. (Currently Amended) The computer-readable storage medium of claim 32,
~~wherein the first subset of the other nodes includes a second node;~~

wherein the method implemented by the program instructions further comprises:

the second node periodically attempting to communicate with each of the other nodes of the plurality of nodes other than the second node;

in response to determining that a particular amount of time has passed without successfully communicating with each of the other nodes of the plurality of nodes other than the second node, the second node initiating an operation to create one or more new replicas of the object.

41. (Currently Amended) The computer-readable storage medium of claim 32,

wherein the first subset of the other nodes includes a second node;

wherein the list of incoherent objects stored on the second node is stored in persistent storage of the second node;

wherein the second node adding the identification of the object to the respective list of incoherent objects stored on the second node comprises:

the second node storing information indicating addition of the identification of the object without modifying the list of incoherent objects stored in the persistent storage of the second node;

after storing the information indicating the addition of the identification of the object, the second node updating the list of incoherent objects stored in the persistent storage of the second node to reflect the addition of the identification of the object.

42. (Previously Presented) The computer-readable storage medium of claim 32,

wherein each respective node updating the replica of the object stored on the respective node comprises each respective node applying a change to the replica of the object stored on the respective node.

43. (Currently Amended) The computer-readable storage medium of claim 32,

wherein the replicas of the object stored on the plurality of nodes includes include a plurality of persistent replicas of the object.

44. (Previously Presented) The computer-readable storage medium of claim 32,
wherein the object is a file;
wherein each replica of the object is a replica of the file.

45. (Currently Amended) The computer-readable storage medium of claim 32,
wherein the method implemented by the program instructions further comprises the first
node receiving an update message;
wherein the first node initiates the ~~communication with each of the other nodes~~ ~~a first~~
~~distributed transaction~~ in response to the update message to attempt to synchronously update the
replicas of the object stored on the other nodes.

46. (Currently Amended) A system comprising:
a plurality of nodes;
wherein the plurality of nodes includes memory storing program instructions executable to
implement a method comprising:
storing a respective replica of an object on each respective node of the plurality of nodes;
a first node of the plurality of nodes initiating ~~communication~~ ~~a first distributed transaction~~
~~with each of the other nodes of the plurality of nodes~~ to attempt to synchronously update the
replicas of the object stored on ~~each of the other nodes of the plurality of nodes~~, wherein ~~the~~
~~communication is successful for each respective node of a first subset of the other nodes~~
~~successfully commits the first distributed transaction to update the replica of the object stored on~~
~~the respective node, and unsuccessful for wherein~~ each node of a second subset of the other nodes
~~does not successfully commit the first distributed transaction; and~~
for each respective node of the first subset of the other nodes:
the respective node updating the replica of the object stored on the respective node;
and
~~each the respective node of the first subset of the other nodes adding an identification of the~~
object to a respective list of incoherent objects stored on the respective node in response to ~~the~~

communication being unsuccessful for the second subset of the other nodes determining that not all of the other nodes successfully committed the first distributed transaction;

wherein the first subset of the other nodes includes a second node of the plurality of nodes;
wherein the method implemented by the program instructions further comprises:

after the second node adds the identification of the object to the respective list of incoherent objects stored on the second node, the second node attempting to communicate with each of the plurality of nodes other than the second node; and

in response to successfully communicating with each of the plurality of nodes other than the second node, the second node initiating a second distributed transaction to synchronize the replicas stored on the second subset of the other nodes with the replicas stored on the first subset of the other nodes.

47. (Currently Amended) The system of claim 46, wherein the method implemented by the program instructions further comprises:

the first node updating the replica of the object stored on the first node; and

the first node adding an identification of the object to a respective list of incoherent objects stored on the first node in response to the communication being unsuccessful for the second subset of the other nodes determining that not all of the other nodes successfully committed the first distributed transaction.

48-49. (Canceled)

50. (Currently Amended) A method comprising:

storing a respective replica of an object on each respective node of a plurality of nodes;
a first node of the plurality of nodes initiating communication a first distributed transaction with each of the other nodes of the plurality of nodes to attempt to synchronously update the replicas of the object stored on each of the other nodes of the plurality of nodes, wherein the communication is successful for each respective node of a first subset of the other nodes successfully commits the first distributed transaction to update the replica of the object stored on

the respective node, and unsuccessful for wherein each node of a second subset of the other nodes does not successfully commit the first distributed transaction; and

for each respective node of the first subset of the other nodes:

the respective node updating the replica of the object stored on the respective node;
and

each the respective node of the first subset of the other nodes adding an identification of the object to a respective list of incoherent objects stored on the respective node in response to the communication being unsuccessful for the second subset of the other nodes determining that not all of the other nodes successfully committed the first distributed transaction;

wherein the first subset of the other nodes includes a second node of the plurality of nodes;
wherein the method further comprises:

after the second node adds the identification of the object to the respective list of incoherent objects stored on the second node, the second node attempting to communicate with each of the plurality of nodes other than the second node; and

in response to successfully communicating with each of the plurality of nodes other than the second node, the second node initiating a second distributed transaction to synchronize the replicas stored on the second subset of the other nodes with the replicas stored on the first subset of the other nodes.

51. (Currently Amended) The method of claim 50, further comprising:
the first node updating the replica of the object stored on the first node; and
the first node adding an identification of the object to a respective list of incoherent objects stored on the first node in response to the communication being unsuccessful for the second subset of the other nodes determining that not all of the other nodes successfully committed the first distributed transaction.

52-53. (Canceled)

Claim 32, 33, 35, 37-47, 50 and 51 (renumbered 1-16) are allowed.

The following is an examiner's statement of reasons for allowance:

None of the prior art of record taken singularly or in combination reasonably teaches or suggests the combined functional limitations of synchronous replication of replicas of objects on plurality of nodes by storing respective replica of an object on every respective node of the plurality of nodes; a first node of the plurality of nodes initiating a first distributed transaction to attempt to synchronously replicate the replica objects stored on each of other (i.e., remaining) nodes of the plurality of nodes; each respective node of a first subset of the other nodes (remaining nodes) successfully committing the first distributed transaction to update the replica of the object stored on the respective node, and wherein each node of a second subset of the other nodes (remaining nodes) does not successfully commit the first distributed transaction; and each respective node of the first subset of the other nodes adding an identification of the object to a respective list of incoherent objects stored on the respective node in response to determining that not all of the other nodes successfully committed the first distributed transaction; wherein the first subset of the other nodes includes a second node of the plurality of nodes; after the second node adds the identification of the object to the respective list of incoherent objects stored on the second node, the second node attempting to communicate with each of the plurality of nodes other than the second node; and in response to successfully communicating with each of the plurality of nodes other than the second node, the second node initiating a second distributed transaction to synchronize the replicas stored on the second subset of the other nodes with the replicas stored on the first subset of the other nodes as recited in claims 32, 46 and 50.

Any comments considered necessary by applicant must be submitted no later than the payment of the issue fee and, to avoid processing delays, should preferably accompany the issue fee. Such submissions should be clearly labeled "Comments on Statement of Reasons for Allowance."

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Yemane Mesfin whose telephone number is (571)272-3927. The examiner can normally be reached on 9:30 AM -7:00 PM.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, William Vaughn can be reached on 572-272-3922. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

/Y. M./
Examiner, Art Unit 2444
